arXiv:1411.5653v2 [stat.ME] 29 Apr 2015

Bayesian Big Data Classification: A Review with Complements

Richard D. Payne, Bani K. Mallick

April 30, 2015

Abstract

This paper focuses on the specific problem of big data classification in a Bayesian setup using Markov Chain Monte Carlo methods. It discusses the challenges presented by the big data problems associated with classification and the existing methods to handle them. Next, a new method based on two-stage Metropolis-Hastings (MH) algorithm is proposed in this context. The purpose of of this algorithm is to reduce the exact likelihood computational cost in the big data context. In the first stage, a new proposal is tested by the approximate likelihood based model. The full likelihood based posterior computation will be conducted only if the proposal passes the first stage screening. Furthermore, this method is adopted into the consensus Monte Carlo framework. Methods are illustrated on two large datasets.

Keywords: Bayesian inference, Big Data, Logistic model, Markov Chain Monte Carlo, Metropolis-Hastings Algorithm

1 Introduction

In the past twenty years, Bayesian statistics have become increasingly popular due to its ability to analyze data of complex structures using flexible models with increased computing power and numerical methods. Consequently, Bayesian methods have been proven to be effective in a wide range of applications. The rise in popularity is largely attributed to simulation based algorithms which can approximate the complex posterior distributions of non-conjugate models, such as Markov Chain Monte Carlo (MCMC) methods including the Metropolis Hastings (MH) algorithm (Roberts and Cassella, 2004).

The term "big data" generally describes data which is too large to be analyzed satisfactorily using current statistical and/or computational methods. Consequently, the actual dimension of "big data" may be quite different depending on the statistical methodology used. When using MCMC methods in the Bayesian paradigm, datasets become "big data" at a much lower dimension than in many frequentist settings due to the repeated computation of the expensive likelihood function. The increased computational time comes from a complete scan of the data through likelihood evaluations on each iteration of the MCMC algorithm.

There are several major methods to overcome this issue. The simplest method involves parallelizing the likelihood to speed up computation. Another method divides the data across multiple machines and performs independent parallel MCMC on each machine to sample from the posterior distribution. The results are then aggregated using weighting (Scott et al., 2013). A third approach is to use subsampling methods to provide a faster estimation of the likelihood (Quiroz et al., 2014; Korattikara et al., 2014; Bardenet et al., 2014).

We propose a method based on a two stage Metropolis algorithm which uses a cheap estimate of the likelihood to determine if a full estimation of the likelihood is necessary. Furthermore, we compare the strengths and weaknesses of the general MCMC Bayesian big data methods of consensus, subsampling, and two-stage Metropolis, as well as investigate the use of a combination of the consensus and two-stage methods. For definiteness, in the following, the focus of this paper is on the classification problem. However, the developed methodology can be extended for any model which is suitable to analyze big data. The methods are applied to a logistic model using individual household loan data from Freddie Mac. Modifications to the techniques described in the papers above have been made to accommodate the features of this dataset and are explained further.

In section 1, we describe the existing methods. In section 2, we develop our proposed method. Section 3 compares the performance of the methods on individual household loan data from Freddie Mac. Section 4 concludes.

2 Model and Methods

We are considering a binary classification problem where the response $\mathbf{y}$ takes the value 0 or 1 and we have covariates $\mathbf{x}$ where $\mathbf{y} = (y_1, \cdots y_n)$ is the vector of responses, $\mathbf{x}$ is the vector of fixed covariates. We use a logit link function to link the ith response with the covariates as

$$\begin{aligned} y_i|\boldsymbol{\beta}, \mathbf{x} &\sim Bernoulli[\pi(\mathbf{x}_i)] \\ \pi(\mathbf{x}_i) &= (1 + e^{-\mathbf{x}_i\boldsymbol{\beta}})^{-1} \\ \boldsymbol{\beta} &\sim MultivariateNormal(\mathbf{0}, \Sigma_0),\ i = 1, \cdots n \end{aligned}$$

where $\boldsymbol{\beta}$ is the l dimensional vector of classification parameters and we assign a Gaussian prior for $\boldsymbol{\beta}$. The model's posterior distribution $p(\boldsymbol{\beta}|\mathbf{y}, \mathbf{x})$ can be expressed as

$$p(\boldsymbol{\beta}|\mathbf{y}, \mathbf{x}) \propto p(\mathbf{y}|\boldsymbol{\beta}, \mathbf{x})p(\boldsymbol{\beta}) \tag{1}$$

where $p(\mathbf{y}|\boldsymbol{\beta}, \mathbf{x})$ is the likelihood function and $p(\boldsymbol{\beta})$ is the prior distribution for the classification parameters. In this non-conjugate prior setup of $\boldsymbol{\beta}$, the posterior distribution $p(\boldsymbol{\beta}|\mathbf{y}, \mathbf{x})$ cannot be expressed in an explicit form and consequently MCMC methods must be used to simulate samples from this posterior distribution. More specifically, we use the Metropolis-Hastings (MH) algorithm to generate samples of $\boldsymbol{\beta}$s from $p(\boldsymbol{\beta}|\mathbf{y}, \mathbf{x})$. The MH algorithm is described as follows.

A.1 MH Algorithm

1. At the tth iteration generate $\boldsymbol{\beta}$ from the proposal distribution $q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)$ where $\boldsymbol{\beta}_t$ is the current state
2. Accept $\boldsymbol{\beta}$ as a posterior sample with probability

$$h(\boldsymbol{\beta}_t, \boldsymbol{\beta}) = min(1, \frac{q(\boldsymbol{\beta}_t|\boldsymbol{\beta})p(\boldsymbol{\beta}|\mathbf{y}, \mathbf{x})}{q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)p(\boldsymbol{\beta}_t|\mathbf{y}, \mathbf{x})}) \tag{2}$$

3. $\boldsymbol{\beta}_{t+1} = \boldsymbol{\beta}$ with probability $h(\boldsymbol{\beta}_t, \boldsymbol{\beta})$ and $\boldsymbol{\beta}_{t+1} = \boldsymbol{\beta}_t$ with probability $1 - h(\boldsymbol{\beta}_t, \boldsymbol{\beta})$.

At each iteration, the probability of moving from the state $\boldsymbol{\beta}_t$ to next state $\boldsymbol{\beta}$ is $q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)h(\boldsymbol{\beta}_t, \boldsymbol{\beta})$, hence the transition kernel for the Markov Chain $\boldsymbol{\beta}_t$ is

$$T(\boldsymbol{\beta}_t, \boldsymbol{\beta}) = q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)h(\boldsymbol{\beta}_t, \boldsymbol{\beta}) + [(1 - \int q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)h(\boldsymbol{\beta}_t, \boldsymbol{\beta})d\boldsymbol{\beta})]I(\boldsymbol{\beta} = \boldsymbol{\beta}_t)$$

where $I()$ is the indicator function. Due to the iterative nature of the algorithm, the likelihood function $p(\mathbf{y}|\boldsymbol{\beta}, \mathbf{x})$ needs to be evaluated repeatedly which is expensive when n is large. Hence, we need to modify the MH algorithm to adapt it for big data problems. In this process, we first describe the existing methods which are suitable for big data classification problems.

2.1 Likelihood Parallelization

Perhaps the simplest way to adapt the Metropolis Hastings algorithm for big data is to compute the likelihood in parallel. In this method, the data are partitioned into p partitions and each assigned to a process/core. On each iteration of the MH algorithm, each core computes the likelihood for its partition and then passes this information (i.e. the sum of the log likelihood) to a central process which aggregates information from each partition. As long as there is no significant communication overhead between the processes, the MH algorithm speed will be increased while still sampling from the true posterior distribution.

2.2 Consensus Monte Carlo

In the consensus Monte Carlo method the data are randomly partitioned into p partitions. Subsequently, allow each partition to run a full MCMC simulation from a posterior distribution given its own data. Lastly, combine the posterior simulations from each partition to produce a set of global draws to reproduce the unified posterior distribution.

Suppose $\mathbf{y} = (y_1, \cdots, y_n)$ denotes the full data and $\mathbf{y}_r$ is the data at the rth partition. We then represent the posterior distribution of $\boldsymbol{\beta}$ as

$$p(\boldsymbol{\beta}|\mathbf{y}, \mathbf{x}) \propto \prod_{j=1}^{p} p(\mathbf{y}_j|\boldsymbol{\beta}) p(\boldsymbol{\beta})^{1/p}$$

where the prior distribution has been expressed as the product of the p components.

For each partition, a Metropolis sampler with a chain of length m is computed in parallel with the prior weight adjusted to p^{-1} its original weight. Once posterior samples are obtained from each of the partitions, the results are combined using a weighted average. The weight, W_i, for the ith partition is equal to the inverse of the posterior covariance matrix obtained from the Metropolis sampler. Let $\boldsymbol{\beta}_i$ be the posterior sample matrix from the ith partition. Thus, the final posterior sample, $\boldsymbol{\beta}$ is obtained using the following weighted average:

$$\boldsymbol{\beta} = \left(\sum_{i=1}^{p} W_i \right)^{-1} \sum_{i=1}^{p} \boldsymbol{\beta}_i W_i.$$

For details see Scott et al. (2013).

2.3 Subsampling based methods

In the subsampling method, a small subset of the data is used to estimate the likelihood function (which is then used to evaluate the acceptance probabilities of the MH algorithm). Since the data size has been reduced, a much faster MCMC algorithm can be developed. From the full data vector $\mathbf{y} = (y_1, \cdots, y_n)$

sample a subset of observations, say $\mathbf{y}_s$. Next evaluate the exact density for this small subset and estimate the likelihood of the complete data, denoted as $\widehat{p}(\mathbf{y}|\boldsymbol{\beta}, X)$, which is usually unbiased. Finally, replace the likelihood $p(\mathbf{y}|\boldsymbol{\beta}, X)$ in the MH algorithm with $\widehat{p}(\mathbf{y}|\boldsymbol{\beta}, X)$. Using the unbiased likelihood estimate in the MCMC chain still provides the correct stationary distribution (Andrieu and Roberts, 2009). Usually complete random sampling does not work well in this situation, but some general guidelines for estimating the full likelihood from a subsample in a Bayesian big data setting have been developed (Quiroz et al., 2014). Two techniques include thin-plate spline approximation and proportional to size sampling using weights created from Gaussian processes. In the logistic classification context, we propose an approximate log-likelihood as in equation 3.

2.4 The Two-Stage Metropolis Hastings algorithm

In the MH algorithm described in A.1, the evaluation of the likelihood is expensive in the big data situation. Generally the MCMC requires thousands of iterations to converge. Furthermore, we need to generate a large number of samples to quantify the uncertainty in the classification coefficients. Additionally, this scheme requires the rejection of several proposals before accepting one from the correct posterior distribution. We use the two-stage MH algorithms where the proposal distribution $q()$ is adapted to the target distribution using an approximate likelihood based model. These algorithms have been used previously (Christen and Fox, 2005; Higdon et al., 2002; Mondal et al., 2014), usually for solving expensive inverse problems. For our purposes, instead of testing each proposal by the exact likelihood based model directly, initially the algorithm tests the proposal by the approximate likelihood based model which is much cheaper to compute. If the proposal is accepted by the initial test, then an exact likelihood based computation will be conducted and the proposal will be further tested as in the MH algorithm method described in A.1. Otherwise, the proposal will be rejected by the approximate model and a new proposal will be generated form $q()$. The approximate likelihood based model filters the unacceptable proposals and avoids the expensive full likelihood computations.

In this particular logit model, we estimate the log-likelihood issuing a variant of the case-control approximate likelihood (Raftery et al., 2012). To understand the approximation, it is important to realize the log-likelihood for a logistic regression model can be written as two sums:

$$
\begin{aligned}
log[p(\mathbf{y}|\boldsymbol{\beta},\mathbf{x})] &= \sum_{i=1}^{n} y_i log[\pi(\mathbf{x_i})] + (1-y_i) log[1-\pi(\mathbf{x_i})] \\
&= \sum_{i=1}^{n} \left\{\theta_i y_i - log(1+e^{\theta_i})\right\} \\
&= \sum_{i:y_i=1} \left\{\theta_i - log(1+e^{\theta_i})\right\} + \sum_{i:y_i=0} -log(1+e^{\theta_i}).
\end{aligned}
$$

where $\theta_i = log[\pi_i/(1-\pi_i)] = \mathbf{x}_i\boldsymbol{\beta}$.

If the data are sparse, then the computation of the first sum will be relatively cheap, and only the second summation needs to be estimated. We used a subsampling method where a random sample of a observations is taken from the failed outcomes (i.e. $y_i = 0$). The second sum in the above equation is estimated by multiplying the average likelihood of the a sampled observations by $n_0 = \sum_{i=1}^{n} I(y_i = 0)$, the number of failures in the dataset. Let A be the index values of the subsample of size a. Thus, the original log-likelihood is estimated as

$$
\widehat{log}[p(\mathbf{y}|\boldsymbol{\beta},\mathbf{x})] = \sum_{i:y_i=1} \left\{\theta_i - log(1+e^{\theta_i})\right\} + \frac{n_0}{a} \sum_{i\in A} -log(1+e^{\theta_i}). \qquad (3)
$$

It is an unbiased estimate of the log-likelihood as $E[\widehat{log}[p(\mathbf{y}|\boldsymbol{\beta},\mathbf{x})] = log[p(\mathbf{y}|\boldsymbol{\beta},\mathbf{x})]$. The likelihood obtained from equation (3) is denoted as $\widehat{p}(\mathbf{y}|\boldsymbol{\beta},\mathbf{x})$. We could obtain an unbiased estimate of the likelihood by making bias correction (Quiroz et al., 2014) but that is not necessary for our method as we are doing further filtering of the proposal using the exact likelihood method.

Using this approximation, the approximate posterior distribution will be

$$
p^*(\boldsymbol{\beta}|\mathbf{y},\mathbf{x}) \propto \widehat{p}(\mathbf{y}|\boldsymbol{\beta},\mathbf{x}) p(\boldsymbol{\beta}).
$$

The two-stage MH sampler proceeds using this estimated log-likelihood $\widehat{p}$ at the first stage.

A.2 Two-Stage MH Algorithm

1. At the tth iteration generate $\boldsymbol{\beta}^{'}$ from the proposal distribution $q(\boldsymbol{\beta}^{'}|\boldsymbol{\beta}_t)$
2. Take a real proposal as

$$
\boldsymbol{\beta} = \begin{cases} \boldsymbol{\beta}^{'} & \text{with probability } \delta(\boldsymbol{\beta}_t, \boldsymbol{\beta}^{'}) \\ \boldsymbol{\beta}_t & \text{with probability } 1-\delta(\boldsymbol{\beta}_t, \boldsymbol{\beta}^{'}) \end{cases}
$$

where

$$
\delta(\boldsymbol{\beta}_t, \boldsymbol{\beta}^{'}) = min\Big(1, \frac{q(\boldsymbol{\beta}_t|\boldsymbol{\beta}^{'}) p^*(\boldsymbol{\beta}^{'}|\mathbf{y},\mathbf{x})}{q(\boldsymbol{\beta}^{'}|\boldsymbol{\beta}_t) p^*(\boldsymbol{\beta}_t|\mathbf{y},\mathbf{x})}\Big)
$$

3. Accept $\boldsymbol{\beta}$ as a posterior sample with probability

$$\rho(\boldsymbol{\beta}_t, \boldsymbol{\beta}) = min\Big(1, \frac{Q(\boldsymbol{\beta}_t|\boldsymbol{\beta})p(\boldsymbol{\beta}|\mathbf{y}, \mathbf{x})}{Q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)p(\boldsymbol{\beta}_t|\mathbf{y}, \mathbf{x})}\Big) \tag{4}$$

where $Q(\boldsymbol{\beta}|\boldsymbol{\beta}_t) = \delta(\boldsymbol{\beta}_t, \boldsymbol{\beta})q(\boldsymbol{\beta}|\boldsymbol{\beta}_t) + (1 - \int \delta(\boldsymbol{\beta}_t, \boldsymbol{\beta})q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)d\boldsymbol{\beta})I(\boldsymbol{\beta} = \boldsymbol{\beta}_t)$

4. Hence take $\boldsymbol{\beta}_{t+1} = \boldsymbol{\beta}$ with probability $\rho(\boldsymbol{\beta}_t, \boldsymbol{\beta})$ and $\boldsymbol{\beta}_{t+1} = \boldsymbol{\beta}_t$ with probability $1 - \rho(\boldsymbol{\beta}_t, \boldsymbol{\beta})$.

At each iteration, the probability of moving from the state $\boldsymbol{\beta}_t$ to next state $\boldsymbol{\beta}$ is $q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)\rho(\boldsymbol{\beta}_t, \boldsymbol{\beta})$, hence the transition kernel for the Markov Chain $\boldsymbol{\beta}_t$ is

$$T(\boldsymbol{\beta}_t, \boldsymbol{\beta}) = q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)\rho(\boldsymbol{\beta}_t, \boldsymbol{\beta}) + (1 - \int q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)\rho(\boldsymbol{\beta}_t, \boldsymbol{\beta})d\boldsymbol{\beta})I(\boldsymbol{\beta} = \boldsymbol{\beta}_t).$$

In the above algorithm, if the trial proposal $\boldsymbol{\beta}^{'}$ is rejected by the approximate posterior then no further computation is needed. Thus, the expensive exact posterior computation can be avoided for those proposals which are unlikely to be accepted. This is just an adaption of the proposal using the approximate posterior where the transition kernel can be written as $K(\boldsymbol{\beta}_t, \boldsymbol{\beta}) = \rho(\boldsymbol{\beta}_t, \boldsymbol{\beta})Q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)$ for $\boldsymbol{\beta} \neq \boldsymbol{\beta}_t$ and $K(\boldsymbol{\beta}_t, \{\boldsymbol{\beta}_t\}) = 1 - \int_{\boldsymbol{\beta} \neq \beta_t} \rho(\boldsymbol{\beta}_t, \boldsymbol{\beta})Q(\boldsymbol{\beta}_t|\boldsymbol{\beta})d\boldsymbol{\beta}$ for $\boldsymbol{\beta} = \boldsymbol{\beta}_t$. This is simple to show that the detailed balance condition $p(\boldsymbol{\beta}_t|\mathbf{y}, \mathbf{x})K(\boldsymbol{\beta}_t, \boldsymbol{\beta}) = p(\boldsymbol{\beta}|\mathbf{y}, \mathbf{x})K(\boldsymbol{\beta}, \boldsymbol{\beta}_t)$ is always satisfied under some minor regularity conditions like the regular MH algorithm.

Result 1: The acceptance probability can be expressed as

$$\rho(\boldsymbol{\beta}_t, \boldsymbol{\beta}) = min\Big(1, \frac{p^*(\boldsymbol{\beta}_t|Y, X)p(\boldsymbol{\beta}|Y, X)}{p^*(\boldsymbol{\beta}|Y, X)p(\boldsymbol{\beta}_t|Y, X)}\Big)$$

Proof: If $\boldsymbol{\beta} = \boldsymbol{\beta}_t$ then the result is trivial since $\rho(\boldsymbol{\beta}_t, \boldsymbol{\beta}) = 1$. For $\boldsymbol{\beta} \neq \boldsymbol{\beta}_t$

$$\begin{aligned}
Q(\boldsymbol{\beta}_t|\boldsymbol{\beta}) &= \delta(\boldsymbol{\beta}, \boldsymbol{\beta}_t)q(\boldsymbol{\beta}_t|\boldsymbol{\beta}) \\
&= \frac{1}{p^*(\boldsymbol{\beta}|Y, X)} min[q(\boldsymbol{\beta}_t|\boldsymbol{\beta})p^*(\boldsymbol{\beta}|Y, X), q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)p^*(\boldsymbol{\beta}_t|Y, X)] \\
&= \frac{q(\boldsymbol{\beta}|\boldsymbol{\beta}_t)p^*(\boldsymbol{\beta}_t|Y, X)}{p^*(\boldsymbol{\beta}|Y, X)}\delta(\boldsymbol{\beta}_t, \boldsymbol{\beta}) \\
&= \frac{p^*(\boldsymbol{\beta}_t|Y, X)}{p^*(\boldsymbol{\beta}|Y, X)}Q(\boldsymbol{\beta}|\beta_t).
\end{aligned}$$

Substituting this in the expression of $\rho(\boldsymbol{\beta}_t, \boldsymbol{\beta})$ we obtain the required expression.

Result 2: When β is promoted from the first stage, then for large a it can be shown that $\rho(\cdot, \cdot)$ goes towards 1. Thus, we only calculate the original full data likelihood when there is a high probability of acceptance of the proposal.

2.5 Combining Consensus and Two-Stage Metropolis

For larger data sets, we propose a combination of the consensus and the two-stage Metropolis methods. This is identical to the consensus method with the

exception that each partition uses the two-stage Metropolis sampler rather than the usual Metropolis sampler. The results are combined using the same weights as the aforementioned consensus method.

3 Data Analysis

The methods introduced above were implemented on two datasets. For initial testing, the methods were used on a relatively small dataset with 50,000 observations from a phone marketing campaign conducted by a Portuguese bank. To test the scaling and speed of the methods, a larger dataset of 16 million observations consisting of individual household loan data from Freddie Mac was used. In both, logistic regression was used to classify observations.

3.1 Analysis of the Portuguese Bank Data

The Portuguese bank dataset was obtained from the University of California Irvine Machine Learning Archive and were analyzed in a recent paper by Moro et al. (2014). The dependent variable of interest was if a client subscribed to a term deposit (binary: yes or no). The predictor variables of interest were the client's previous promotion outcome (non-existent, failure, success), age (years), and type of contact (telephone, cellular). The model was parameterized as follows:

$$log\left(\frac{\pi(\mathbf{x}_i)}{1-\pi(\mathbf{x}_i)}\right) = \beta_0 + \beta_1 failure + \beta_2 success + \beta_3 age + \beta_4 cellular$$

where failure, success, and cellular are indicator variables.

For both, a logistic model was used with a multivariate normal prior on $\boldsymbol{\beta} = \{\beta_0, \beta_1, \beta_2, \beta_3, \beta_4\}^T$ and a logit likelihood:

$$\begin{aligned} y_i|\boldsymbol{\beta}, \mathbf{x} &\sim Bernoulli[\pi(\mathbf{x}_i)] \\ \pi(\mathbf{x}_i) &= (1 + e^{-\mathbf{x}_i\boldsymbol{\beta}})^{-1} \\ \boldsymbol{\beta} &\sim MultivariateNormal(\mathbf{0}, \Sigma_0 = I * 1000),\ i = 1, \cdots n \end{aligned}$$

where Σ_0 is the covariance matrix for $\boldsymbol{\beta}$, I is the identity matrix and $\mathbf{x}_i$ is the ith row of the design matrix.

For each method, a Metropolis chain was produced with 500,000 samples, burn-in of 1000 and thinning every 20 draws. Figures 1-3 compare the posterior densities of the consensus, subsampling, and two-stage methods to the usual Metropolis sampler results.

Based on estimated densities, it appears the two-stage Metropolis is most accurate, followed by the consensus method and subsampling techniques. We

Figure 1: Posterior densities from the Portuguese bank data, consensus Monte Carlo vs. Metropolis-Hastings.

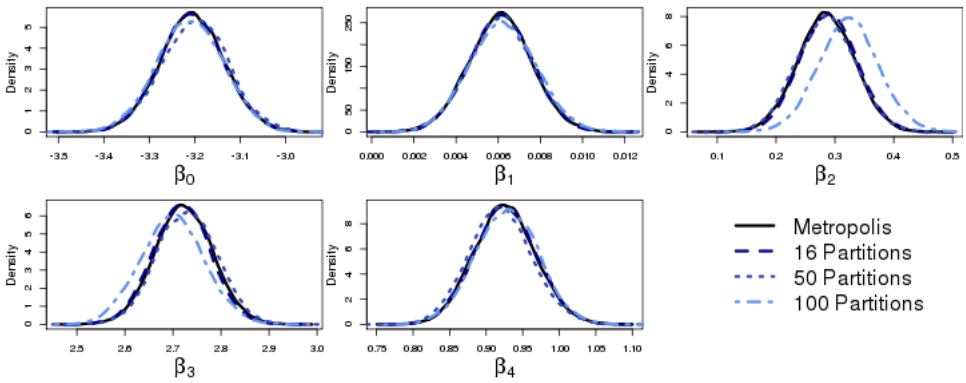

Figure 2: Posterior densities from the Portuguese bank data, subsampling vs. Metropolis-Hastings.

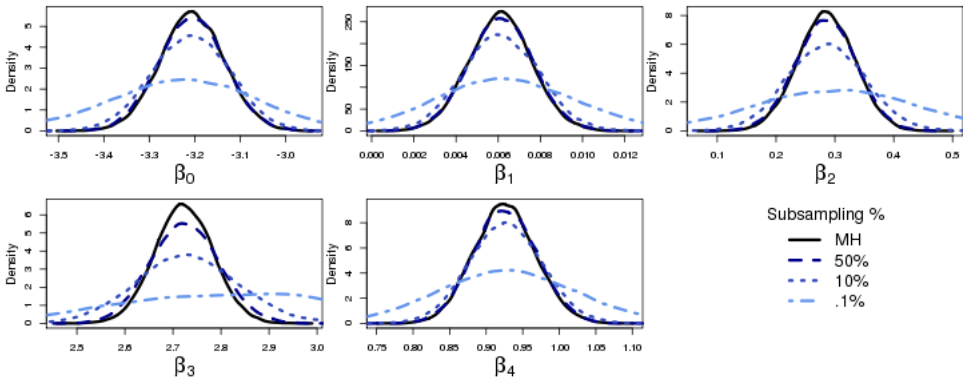

Figure 3: Posterior densities from the Portuguese bank data, two-stage Metropolis vs. Metropolis-Hastings.

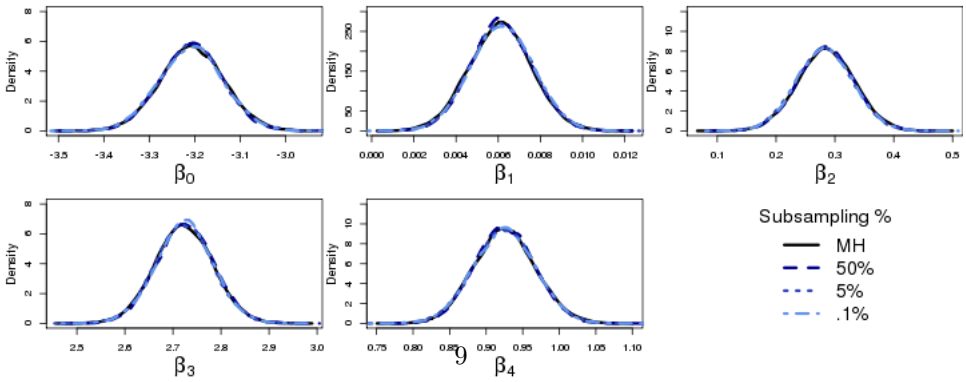

note the two-stage Metropolis method was very accurate even when subsampling only .1% of the data. Furthermore, it requires no bias corrections which may be necessary for the consensus method when partitions are not sufficiently large (Scott et al., 2013). For this dataset, the consensus method produced biased posteriors when there were a large number of partitions (jackknife bias correction can help in some cases but was not employed here (Scott et al., 2013)). The subsampling method produced posterior intervals which increased in width as subsampling percentage decreased. Quiroz et al. (2014) indicate this is due to the random sampling of the data and suggest using fairly sophisticated probability proportional-to-size sampling techniques which were not employed here. Although not plotted, the combination of the consensus and two-stage methods produced results that were identical to the consensus method within Monte Carlo error.

Since the Portuguese bank datset is relatively small (about 50,000 observations), we refer readers to the next section for a careful comparison of the speeds of the different methods on the much larger Freddie Mac dataset (16 million observations).

3.2 Analysis of Freddie Mac Data

The larger Freddie Mac dataset (16 million observations) was used to further test these methods and provide speed comparisons between the above methods. With the exception of the subsampling method, the methods introduced above were implemented on household loan data from from Freddie Mac. All methods were computed on a simulation server with 16 cores and the likelihood was parallelized to provide a level playing field in terms of speed. A logistic model was used with a multivariate normal prior on $\boldsymbol{\beta}$ and a logit likelihood:

$$\begin{aligned} y_i|\boldsymbol{\beta}, \mathbf{x} &\sim Bernoulli[\pi(\mathbf{x}_i)] \\ \pi(\mathbf{x}_i) &= (1 + e^{-\mathbf{x}_i\boldsymbol{\beta}})^{-1} \\ \boldsymbol{\beta} &\sim MultivariateNormal_l(\mathbf{0}, \Sigma_0 = I * 1000),\ i = 1, \cdots n \end{aligned}$$

where Σ_0 is the covariance matrix for $\boldsymbol{\beta}$, I is the identity matrix and $\mathbf{x}_i$ is the ith row of the design matrix.

The loan data from Freddie Mac were obtained in February 2015 from Freddie Mac's website. The data consists of nearly 16 million loans which Freddie Mac acquired during 1999-2012 and contains monthly performance data on each loan. The dependent variable of interest is whether or not a loan was at any point 30 or more days delinquent within the first year of origination. For each year, a logistic model with intercept was fitted with FICO scores, first-time homebuyer status (no, yes, unknown), debt to income ratio, original unpaid balance, and loan purpose (purchase, cash-out refinance, no cash-out refinance)

Table 1: Speed Comparison

Method	**Iterations per core per second**	**Speed Ratio**
Consensus & Two-Stage	20.2	1.92
Two-Stage	19.7	1.87
Consensus	12.0	1.14
Parallel	10.5	–

as variables of interest ($l = 112$ total parameters). By fitting the model for each year, we can observe how each variable changed before, during, and after the great economic recession.

For each method, a parallel (16 cores) Metropolis chain was produced with 500,000 iterations and burn-in of 50,000. The dependent variable is fairly sparse – approximately 4.6% of the loans had any delinquency within the first year. In the two-stage Metropolis algorithm, all delinquent loans and a sample of 16,000 non-delinquent loans were used to estimate the log-likelihood. For brevity, we examine only the posterior density plots of fico credit score for the years 1999-2012 (similar results were obtained for other predictors). Figures 4-5 compare the results of the consensus and two-stage methods to the parallelized MH algorithm as a reference.

The plots indicate that both methods performed fairly well on this dataset. Since both the parallelized likelihood and the two-stage Metropolis algorithm converge to the true posterior distribution, any differences can be attributed to Monte Carlo error. The consensus method performed very well with the exception of slight bias in a few of the posterior distributions. In a smaller test set of this data, increasing the number of partitions in the consensus method increased the bias in the posterior intervals which is consistent with Scott et al. (2013). However, due to the large number of observations in each partition (nearly 1 million), this bias is arguably negligible for this application. Although not pictured, the results of the combined two-stage and consensus method were identical to the consensus method within Monte Carlo error.

Table 1 compares the speeds of the different methods. The methods involving two-stage Metropolis performed fastest. In a pilot analysis using a subset of the data (100,000 loans), there was a general increase in speed as subsampling percentage decreased for the two-stage method. However, for extremely small subsampling percentages the two-stage method began to slow. This is likely due to a less accurate approximation of the likelihood in the first stage leading to unnecessary evaluations of the full likelihood function. Further work needs to be done to estimate an optimal subsampling percentage.

Figure 4: Posterior densities from Freddie Mac data, consensus Monte Carlo vs. parallel Metropolis-Hastings.

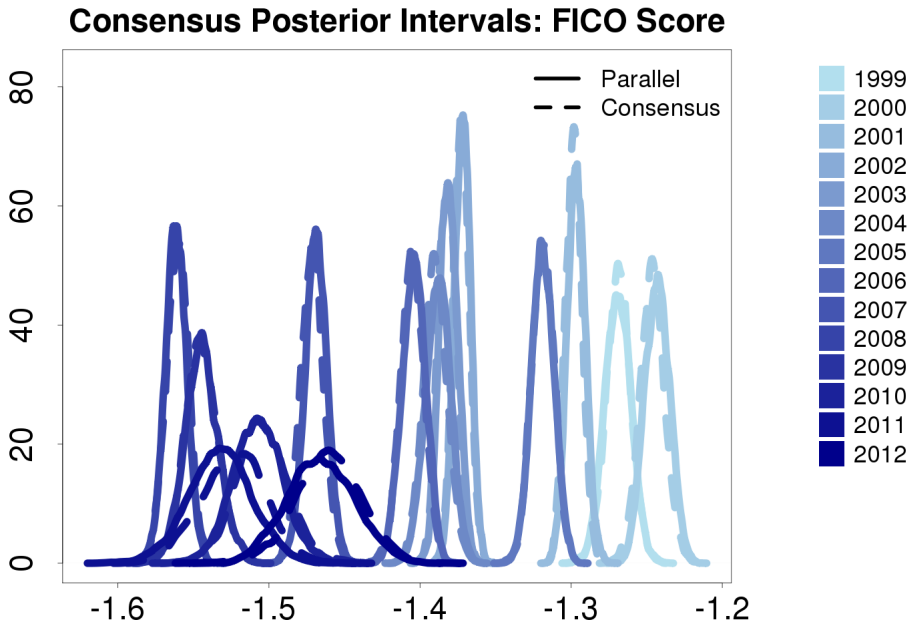

Figure 5: Posterior densities from Freddie Mac data, two-stage Metropolis vs. parallel Metropolis-Hastings.

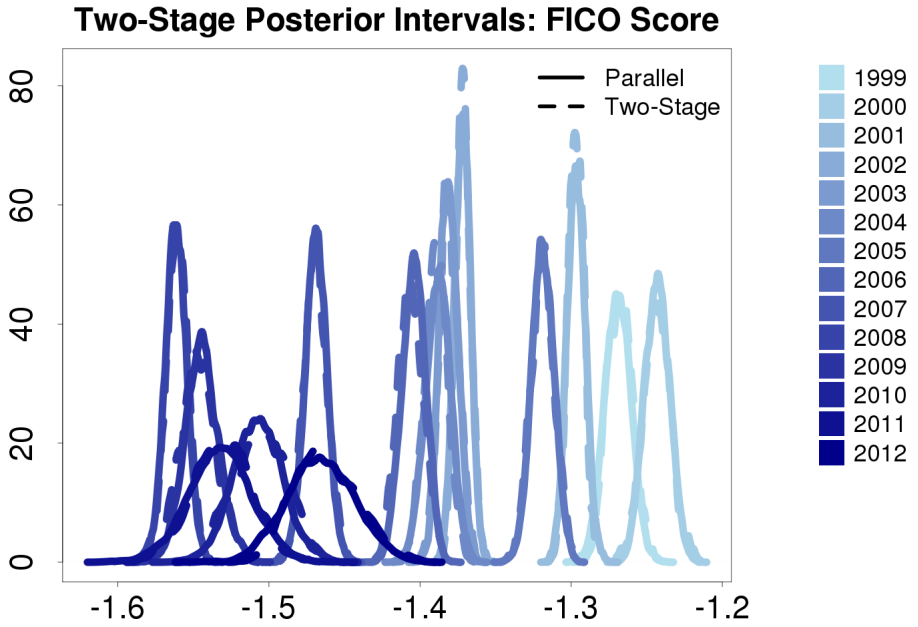

4 Conclusion

The results from this analysis indicate there are a number of big data Bayesian methods which are effective in obtaining/approximating the posterior distribution more quickly than traditional methods. The two-stage Metropolis appears to be a very fast and scalable method since it is accurate even when a small percentage of the data is used. The introduction of a two-stage Metropolis combined with the consensus method shows promise for even larger datasets.

It is important to note there exists a trade-off between the accuracy and the scalability of the methods above. Methods which sample from the true posterior are typically limited to datasets which can fit within RAM (i.e. parallel likelihood, two-stage method) whereas methods which can handle larger datasets that may not fit in RAM have no theoretical guarantees for convergence to the true posterior distribution (i.e. consensus and combined consensus/two-stage methods). Further work needs to be done to develop methods which are scalable to datasets which are too large to be contained within RAM and also converge to the true posterior distribution.

Other extensions to this work include finding ways to randomly select partitions for use in the consensus method and combining existing and new methods in other innovative ways. Furthermore, the proposed algorithms will be generalized to handle data adaptive nonlinear classification problems.

References

Andrieu, C. and Roberts, G. O. (2009). "The pseudo-marginal approach for efficient Monte Carlo computations." *Ann. Statist.*, 37(2): 697–725. URL `http://dx.doi.org/10.1214/07-AOS574`

Bardenet, R., Doucet, A., and Holmes, C. (2014). "Towards scaling up Markov Chain Monte Carlo: an adaptive subsampling approach." *Technical report*.

Christen, J. and Fox, C. (2005). "Markov chain Monta Carlo using an approximation." *Journal of Computational and Graphical Statistics*, 21:4.

Higdon, D., Lee, H., and Bi, Z. (2002). "A Bayesian approach to characterizing uncertainty in inverse problems using coarse and fine scale information." *IEEE Trans. Signal Processing*, 22.

Korattikara, A., Chen, Y., and Welling, M. (2014). "Austerity in MCMC land: Cutting the Metropolis-Hastings budget." *Proceedings of the International Conference on Machine learning*.

Mondal, A., Efendiev, Y., Dattagupta, A., and Mallick, B. (2014). "Bayesian unceratinty quantification for subsurface inversion using mutiscale hierarchical model." *Technometrics*, 56.

Moro, S., Cortez, P., and Rita, P. (2014). "A Data-Driven Approach to Predict the Sucess of Bank Telemarketing." *Decision Support Systems, Elsevier*, 62: 22–31.

Quiroz, M., Villani, M., and Kohn, R. (2014). "Speeding Up MCMC by Efficient Data Subsampling." *ArXiv e-prints*.

Raftery, A. E., Niu, X., Hoff, P. D., and Yeung, K. Y. (2012). "Fast Inference for the Latent Space network Model Using a Case-Control Approximate Likelihood." *Journal of Computational and Graphical Statistics*, 21:4: 901–919.

Roberts, C. and Cassella, G. (2004). "Monte Carlo Statistical Methods." *Springer-Verlag*.

Scott, S. L., Blocker, A. W., and Bonassi, F. V. (2013). "Bayes and Big Data: The Consensus Monte Carlo Algorithm."